

Maximum *a Posteriori* State Path Estimation: Discretization Limits and their Interpretation [★]

Dimas Abreu Dutra, Bruno Otávio Soares Teixeira, Luis Antonio Aguirre.

Programa de Pós-Graduação em Engenharia Elétrica - Universidade Federal de Minas Gerais - Av. Antônio Carlos 6627, 31270-901, Belo Horizonte, MG, Brasil

Abstract

Continuous-discrete models with dynamics described by stochastic differential equations are used in a wide variety of applications. For these systems, the maximum *a posteriori* (MAP) state path can be defined as the curves around which lie the infinitesimal tubes with greatest posterior probability, which can be found by maximizing a merit function built upon the Onsager–Machlup functional. A common approach used in the engineering literature to obtain the MAP state path is to discretize the dynamics and obtain the MAP state path for the discretized system. In this paper, we prove that if the trapezoidal scheme is used for discretization, then the discretized MAP state path estimation converges hypographically to the continuous-discrete MAP state path estimation as the discretization gets finer. However, if the stochastic Euler scheme is used instead, then the discretized estimation converges to the minimum energy estimation. The minimum energy estimates are, in turn, proved to be the state paths associated with the MAP *noise* paths, which in some cases differ from the MAP *state* paths. Therefore, the discretized MAP state paths can have different interpretations depending on the discretization scheme used.

Key words: Estimation theory; optimization problems; variational analysis; stochastic systems.

1 Introduction

A wide variety of phenomena of engineering interest are continuous-time in nature and can be modelled by stochastic differential equations (SDEs). In practical situations, the measurements of these systems are usually sampled, leading to continuous-discrete models, i.e., models with continuous-time dynamics and measurements taken in discrete-time. For this class of models, the maximum *a posteriori* (MAP) state paths can be defined as the curves which have the greatest posterior sojourn probability, i.e., that maximize the probability that the state process lies inside an infinitesimal tube around them. The asymptotic prior probability of ϵ -tubes around state paths, as ϵ vanishes, is given by the Onsager–Machlup functional [Fujita and Kotani, 1982, Zeitouni and Dembo, 1987, Capitaine, 1995], which can be used to construct a variational problem for MAP state path estimation in continuous-discrete models [Aihara and Bagchi, 1999b]. It should be noted, however,

that variational MAP state path estimation using the Onsager–Machlup functional is not common in the engineering literature. One of the few exceptions is the work of Aihara and Bagchi [1999a,b], where only experiments with simulated data are shown.

In the context of discrete-time dynamical systems, there exist several methods for maximum *a posteriori* (MAP) state path estimation [Bell, 1994, Godsill et al., 2001, Aravkin et al., 2011a, Monin, 2013]. The MAP state path consists of the joint posterior mode of the states over all time indices, given a series of measurements. These estimators have been the focus of many recent publications due to their robustness properties and larger applicability than nonlinear Kalman smoothers [Farahmand et al., 2011, Aravkin et al., 2011a,b, 2012a,b, Dutra et al., 2012, Monin, 2013].

To use discrete-time MAP state path estimation in the context of continuous-discrete systems, the system dynamics needs to be converted to discrete-time. However, while the discretization of linear systems can be exact, for general nonlinear SDEs it is necessary to employ approximations such as the stochastic Euler, trapezoidal or Itô–Taylor schemes [Kloeden and Platen, 1992]. Applications of MAP state path estimation on discretized

[★] This work has been financially supported by the Brazilian agencies CAPES, CNPq, FAPEMIG and FINEP.

Email addresses: ddutra@cpdee.ufmg.br (Dimas Abreu Dutra), brunooot@ufmg.br (Bruno Otávio Soares Teixeira), aguirre@cpdee.ufmg.br (Luis Antonio Aguirre).

continuous-discrete systems described by SDEs are presented in the work of Bell et al. [2009] and Aravkin et al. [2011a, 2012a]. In these applications, the former used the exact discretization for linear systems, while the latter used the Euler discretization for nonlinear systems.

Discrete-time approximations of SDEs improve as the discretization step decreases. Applications in Kalman filtering that use such discretization schemes, for example, usually require fine discretizations to produce meaningful results [Särkkä and Solin, 2012], especially when the measurements are sparse. An important question then arises: how are the MAP state paths of the discretized systems related to the MAP state path of the original system? Do the discretized estimates converge to continuous-time functions?

In this paper, we use the theory of consistent discretizations [Polak, 1997] to show that, under some regularity conditions, the MAP estimates using the Euler and trapezoidal schemes converge to maxima of variational problems. The limit variational problem for the trapezoidal discretization is the MAP state path estimation problem. For the Euler discretization, however, the limit variational problem is the minimum energy estimation problem [Cox, 1963, Sec. 3.7], whose merit function is similar to that of the MAP problem but lacks a correction term which is related to the amplification of noise by the drift. The minimum energy estimator was, for many years, believed to be a good MAP estimator [Zeitouni and Dembo, 1987, p. 234]. Its estimates are MAP state paths, however, only when the correction term is constant. In this paper we prove that, in the general case, the minimum energy estimates correspond to the state paths associated with the MAP *noise* paths.

The main contributions of this paper are *i)* providing the statistical interpretation for minimum energy estimates as the state path associated with the MAP noise path, *ii)* proving a link between variational and discretized MAP state path estimation, and *iii)* proving that the limits of the discretized MAP state path estimation problems depend on the choice of discretiza-




Fig. 1. Comparison of the different MAP estimators for a simple SDE. The discretizations with the exact transition densities lie on the state path obtained with the Onsager–Machlup, for all N , and are not shown to avoid clutter.

tion used. The remainder of this paper is organized as follows. In Sec. 2 a simple motivating example is presented to illustrate the main ideas of the paper. In Sec. 3 we present a formal definition of mode, in the framework of Bayesian decision theory, that generalizes well to infinite-dimensional spaces. We then define the continuous-discrete estimation problems in Secs. 4–6 and the discretized problems in Sec. 7. Finally, in Sec. 8 the hypo-convergence of the problems is proved. Simulated examples are presented in Sec. 9 and the conclusions and future work are presented in Sec. 10.

2 Motivating example

Take the simple scalar SDE below, for which the exact transition probability is known [Daum, 1986]:

$$dX_t = \tanh(X_t) dt + dW_t.$$

Then, consider a state path estimation problem where the only information available is the initial state prior distribution and a single measurement at $t = 5$ with value $y = 1.5$. The initial state and measurement are normally distributed with mean and covariance as follows:

$$X_0 \sim \mathcal{N}(0, 0.16), \quad Y|X_5 \sim \mathcal{N}(X_5, 0.16).$$

If we discretize this SDE at N equally spaced time-points using the exact transition density, the Euler scheme and the trapezoidal scheme, it can be seen that the MAP state path estimates converge to continuous-time functions as $N \rightarrow \infty$. The estimates with the trapezoidal and exact transition densities have the same limit, which is the MAP state path obtained using the Onsager–Machlup functional. *The estimates with the Euler discretization, however, converge to a different curve, which is the minimum energy state path.* These results are summarized in Fig. 1. The maximization problems corresponding to each solution are presented in Secs. 6–7.

This example shows that, for some systems, the Euler discretization is not appropriate for approximating the joint density of the state path, even with very fine discretizations. The MAP state paths obtained with the Euler scheme converge to the minimum energy estimates, not to the true MAP state paths obtained with the exact transition density or the Onsager–Machlup functional.

3 Bayesian decision and MAP estimation

This section is a short introduction to Bayesian estimation theory to argue that the maximization of the probability of tubes of vanishing radius is a natural extension of the concept of MAP estimation to continuous-time stochastic processes. We also formalize the concept with

a clear mathematical definition that takes multimodality into account.

To begin, let (Ω, \mathcal{E}, P) be the probability space on which all random variables are defined. Random variables will be denoted by uppercase letters and their values by lowercase, so that if $X: \Omega \rightarrow \mathcal{X}$ is a \mathcal{X} -valued random variable, $x \in \mathcal{X}$ will be used to refer to specific values it might take. The same applies for stochastic processes. The dependency on the random outcome $\omega \in \Omega$ will be omitted when unambiguous, to simplify the notation.

Bayesian estimation can be formulated as taking a decision which minimizes an expected loss, given an observation. Say X is a random variable, which is to be estimated, and \mathcal{A} is an observed event, summarizing all the information available. The Bayesian choice [Robert, 2001] for the estimate \hat{x} corresponds to a value which minimizes of the expected loss, i.e.,

$$\mathbb{E}[\ell(X, \hat{x}) | \mathcal{A}] = \inf_x \mathbb{E}[\ell(X, x) | \mathcal{A}],$$

where ℓ is an \mathbb{R} -valued function which represents the loss associated with each outcome of X and choice for x .

Several known loss functions correspond to certain statistics of the posterior distribution. When the loss is the absolute distance and squared distance between the estimate and the outcome of the random variable,

$$\ell_1(X, x) := \|X - x\|, \quad \ell_2(X, x) := \|X - x\|^2,$$

then the Bayesian estimator is the posterior median and mean, respectively. For discrete random variables, i.e., when the image of X is a countable set, the posterior modes are the Bayesian estimates associated with the 0–1 loss function, given by

$$\ell_{01}(X, x) := 0 \text{ if } X = x, 1 \text{ otherwise.}$$

For \mathbb{R}^n -valued random variables which admit continuous posterior densities, the expected value of ℓ_{01} is always unity, so we must instead consider a family of 0–1 loss functions given by

$$\ell_\epsilon(X, x) := 0 \text{ if } \|X - x\| \leq \epsilon, 1 \text{ otherwise.}$$

One possible definition for the MAP estimate is then the limit of Bayesian estimates \hat{x}_ϵ associated with the losses ℓ_ϵ , as $\epsilon \downarrow 0$ [Robert, 2001, Sec. 4.1.2]. It can be shown that the limit of any convergent subsequence of \hat{x}_ϵ is always a maximum of the posterior density.

However, the above definition for MAP estimate has several problems. First, convergence of the Bayesian estimates is hard to guarantee. Even the existence of convergent subsequences requires additional assumptions on the distribution of the random variable. In addition,

there can exist maxima of posterior density that cannot be limits of the Bayesian estimates. These limitations are often overlooked because the *de facto* definition of MAP estimates is the location of the maxima of the posterior density, i.e., the modes of the posterior distribution. For random variables in infinite-dimensional spaces, for which there is no probability density in the usual sense, this *de facto* definition cannot be used. To overcome the problems mentioned above, we present below a formal definition of mode which agrees with the *de facto* one and generalizes well to infinite-dimensional spaces.

Definition 1 (mode). Let \mathcal{X} be a normed vector space and X be a \mathcal{X} -valued random variable. An element $\hat{x} \in \mathcal{X}$ is a mode of X if, for all $x \in \mathcal{X}$,

$$\limsup_{\epsilon \downarrow 0} \frac{P(\|X - x\| \leq \epsilon)}{P(\|X - \hat{x}\| \leq \epsilon)} \leq 1.$$

Definition 2. The maximum *a posteriori* (MAP) estimates of a random variable are the modes of its posterior distribution, according to Definition 1.

Definitions 1 and 2 can be interpreted in the framework of Bayesian decision theory because

$$\mathbb{E}[\ell_\epsilon(X, x) | \mathcal{A}] = 1 - P(\|X - x\| \leq \epsilon | \mathcal{A}).$$

Therefore, if \hat{x} is a MAP estimate and x is not, there exists an $\varepsilon \in \mathbb{R}_{>0}$ such that, for all $\epsilon < \varepsilon$, the expected ℓ_ϵ loss associated with \hat{x} is no greater than that associated with x . If two points are MAP estimates, then the ratio of the expected ℓ_ϵ losses associated with them converges to unity. Furthermore, if X is an \mathbb{R}^n -valued random variable that admits a continuous density p_X , then for all $a, b \in \mathbb{R}^n$ with $p_X(b) > 0$

$$\lim_{\epsilon \downarrow 0} \frac{P(\|X - a\| \leq \epsilon)}{P(\|X - b\| \leq \epsilon)} = \frac{p_X(a)}{p_X(b)}, \quad (1)$$

so a point $x \in \mathbb{R}^n$ is a MAP estimate if and only if it maximizes the density. Consequently, Definition 1 agrees with the *de facto* definition of mode.

Having formalized the definition of mode for general random variables, we can now show how the Onsager–Machlup functional can be used to obtain the mode of the state paths of systems described by SDEs.

4 Probability of tubes around state paths

Let X_t be an \mathbb{R}^n -valued stochastic process representing the state of a dynamic system at the instant t over the time interval $\mathcal{T} := [0, T]$. The evolution of X is given by the SDE

$$dX_t = f(t, X_t) dt + G dW_t, \quad (2)$$

where $f: \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the drift function, $G \in \mathbb{R}^{n \times n}$ is the diffusion matrix and W_t is an n -dimensional Wiener process which represents the process noise.

We now make some basic assumptions about the state process and the functions which define its dynamics. A function will be said to be of class C_b^n if it is n times differentiable and all its derivatives up to degree n , including the zeroth, are bounded and continuous.

Assumption 3 (state process).

- a. The drift f is continuous with respect to both arguments and is of class C_b^2 with respect to its second argument x , uniformly over its first argument t .
- b. The diffusion matrix G has full rank.
- c. The initial state X_0 is independent of the Wiener process W .
- d. The initial state X_0 admits a continuous density $\nu(x_0)$ of class C_b^0 .

Together, these assumptions allow the calculation of the asymptotic probability of ϵ -tubes of the state process, also called ϵ -sausages. The ϵ -tube around a given $\phi: \mathcal{T} \rightarrow \mathbb{R}^n$ corresponds to all curves which lie within a distance ϵ of ϕ , with respect to the supremum norm. We denote by $\mathcal{B}_\phi^\epsilon \in \mathcal{E}$ the event that the state process is contained within the ϵ -tube around ϕ , i.e.,

$$\mathcal{B}_\phi^\epsilon := \{\omega \in \Omega: \sup_{t \in \mathcal{T}} \|X_t(\omega) - \phi(t)\| \leq \epsilon\},$$

where, for elements of \mathbb{R}^n , $\|\cdot\|$ denotes the Euclidean norm. The probability of $\mathcal{B}_\phi^\epsilon$ is also referred to as the sojourn probability around ϕ . The tube events define regions of the $\mathcal{T} \times \mathbb{R}^n$ space, as illustrated in Fig. 2, and an outcome $\omega \in \Omega$ belongs to the tube event if all of its pairs $(t, X_t(\omega))$ lie in the shaded region.

Of special interest are the tubes around absolutely continuous functions with square-integrable derivatives, i.e., the Cameron–Martin space \mathcal{H} of the law of the Wiener process:

$$\mathcal{H} := \{(x: \mathcal{T} \rightarrow \mathbb{R}^n): x \text{ abs. continuous, } \dot{x} \in L^2(\mathcal{T})\}.$$

For any such path $\phi \in \mathcal{H}$, the Onsager–Machlup functional is defined as

$$J(\phi) := -\frac{1}{2} \int_0^T [\|\dot{\phi}(t) - f(t, \phi(t))\|_Q^2 + \text{div } f(t, \phi(t))] dt, \quad (3)$$

where $Q := (GG^\top)^{-1}$, $\|x\|_A := \sqrt{x^\top A x}$ and $\text{div } f$ is the divergence of the drift vector field. The importance of this functional is outlined in the theorem below.

Theorem 4 (Fujita and Kotani [1982]). *Consider a system (2) with state process satisfying Assumption 3. Then,*




Fig. 2. Illustration of the tubes corresponding to $\mathcal{B}_\phi^\epsilon$ and $\mathcal{B}_\phi^\epsilon$.

for all $\phi, \varphi \in \mathcal{H}$ with $\nu(\varphi(0)) > 0$,

$$\lim_{\epsilon \downarrow 0} \frac{P(\mathcal{B}_\phi^\epsilon)}{P(\mathcal{B}_\varphi^\epsilon)} = \frac{\exp(J(\phi)) \nu(\phi(0))}{\exp(J(\varphi)) \nu(\varphi(0))}. \quad (4)$$

where J is the Onsager–Machlup functional defined in (3) and ν is the initial state density satisfying Assumption 3d.

Comparing (4) to (1), Thm. 4 implies that, for the purpose of obtaining the mode, the product of the exponential of the Onsager–Machlup functional and the initial state density plays a role analogous to a probability density. For this reason, these terms are sometimes referred to as a *fictitious* density of state paths [Zeitouni, 1989]. Also note that (4) is valid for tubes with respect to other norms of function spaces besides the supremum norm, and that the limit does not depend on the norm used [Capitaine, 1995].

For tubes around $\gamma \notin \mathcal{H}$, outside the Cameron–Martin space, we assume that their probability decays much faster than that of tubes around $\varphi \in \mathcal{H}$ that belong to the Cameron–Martin space, i.e., $\lim_{\epsilon \downarrow 0} P(\mathcal{B}_\gamma^\epsilon)/P(\mathcal{B}_\varphi^\epsilon) = 0$. This implies that the mode of the state path, according to Def. 1, lies in \mathcal{H} and can be obtained by maximizing the denominator of the right-hand side of (4). This assumption is implicit whenever the Onsager–Machlup functional is used to obtain the mode of state paths [Zeitouni and Dembo, 1987, Aihara and Bagchi, 1999a,b].

Whenever the divergence of the drift vector field does not depend on the state x , then the last term of the integral in (3) can be dropped from the Onsager–Machlup functional, leading to the energy functional. Even for systems with state-dependent drift divergence, the energy functional gives the asymptotic probability of ϵ -tubes around the path of the Wiener process associated with the state path, as is shown in the next section.

5 Probability of tubes around noise paths

The Wiener process is also a diffusion process for which the Onsager–Machlup functional can be used for calculating the asymptotic probability of ϵ -tubes. For any $\eta: \mathcal{T} \rightarrow \mathbb{R}^n$ with $\eta(0) = 0$ and $x_0 \in \mathbb{R}^n$, let $\mathcal{C}_{\eta, x_0}^\epsilon \in \mathcal{E}$ represent the event that the initial state and the Wiener

process lie within an ϵ radius of x_0 and η , i.e.,

$$\mathcal{C}_{\eta, x_0}^\epsilon := \{\omega \in \Omega: \sup_{t \in \mathcal{T}} \|W_t(\omega) - \eta(t)\| \leq \epsilon, \\ \|X_0(\omega) - x_0\| \leq \epsilon\}.$$

Applying a variant of Thm. 4 which considers diffusions with a fixed initial condition [cf. Fujita and Kotani, 1982], we have that, for any $\zeta, \eta \in \mathcal{H}$ and $\chi_0, x_0 \in \mathbb{R}^n$ such that $\zeta(0) = \eta(0) = 0$ and the initial state density $\nu(x_0) > 0$,

$$\lim_{\epsilon \downarrow 0} \frac{P(\mathcal{C}_{\zeta, \chi_0}^\epsilon)}{P(\mathcal{C}_{\eta, x_0}^\epsilon)} = \frac{\exp\left(-\frac{1}{2} \int_0^T \dot{\zeta}(t)^\top \dot{\zeta}(t) dt\right) \nu(\chi_0)}{\exp\left(-\frac{1}{2} \int_0^T \dot{\eta}(t)^\top \dot{\eta}(t) dt\right) \nu(x_0)}. \quad (5)$$

To relate the sojourn probability of state paths and noise paths, we now show that the event corresponding to each noise path tube is contained in an event corresponding to a state path tube.

Lemma 5. *If the system satisfies Assumption 3, then there exists a constant $c_1 \in \mathbb{R}$ such that $\mathcal{C}_{\eta, \phi(0)}^\epsilon \subset \mathcal{B}_\phi^{c_1 \epsilon}$ for all $\epsilon \in \mathbb{R}$ and $\eta, \phi \in \mathcal{H}$ such that $\eta(0) = 0$ and*

$$\dot{\phi}(t) = f(t, \phi(t)) + G\dot{\eta}(t). \quad (6)$$

Proof. The SDE (2) can be written in the integral form as

$$X_t = X_0 + \int_0^t f(t, X_t) dt + GW_t.$$

Writing the associated ordinary differential equation (ODE) (6) in the same form and taking the difference, we have that

$$\|X_t - \phi(t)\| \leq \|X_0 - \phi(0)\| + \|GW_t - G\eta(t)\| \\ + \int_0^t \|f(t, X_t) - f(t, \phi(t))\| dt.$$

Since the Jacobian matrix of f is assumed to be bounded, f is Lipschitz continuous, i.e., there exist $K \in \mathbb{R}_{>0}$ such that

$$\|f(t, a) - f(t, b)\| \leq K \|a - b\| \quad \forall t \in \mathcal{T} \text{ and } a, b \in \mathbb{R}^n.$$

Defining $c_2 := 1 + \|G\|$, where the norm of G is induced by the Euclidean norm, we have that, for all $\omega \in \mathcal{C}_{\eta, \phi(0)}^\epsilon$,

$$\|X_t - \phi(t)\| \leq c_2 \epsilon + \int_0^t K \|X_t - \phi(t)\| dt.$$

Using the Bellman–Grönwall inequality we then have that

$$\|X_t - \phi(t)\| \leq \epsilon c_2 \exp(TK). \quad \square$$

Lem. 5 means that every state path in \mathcal{H} is associated with a noise path in \mathcal{H} and a initial state through the associated ODE (6). Rewriting (5) in terms of the state path ϕ associated with the noise path η , we note that

$$-\frac{1}{2} \dot{\eta}(t)^\top \dot{\eta}(t) = -\frac{1}{2} \|\dot{\phi}(t) - f(t, \phi(t))\|_Q^2,$$

which leads to the energy functional:

$$J_e(\phi) := -\frac{1}{2} \int_0^T \|\dot{\phi}(t) - f(t, \phi(t))\|_Q^2 dt, \quad (7)$$

so called because it corresponds to the energy of the noise signal associated with each state path. The energy functional is the Onsager–Machlup functional without the drift divergence term, which accounts for the amplification of noise by the system dynamics.

Next, we show how to incorporate the measurement likelihood so that the asymptotic *posterior* sojourn probability can be obtained and, with it, the MAP estimator.

6 Continuous-discrete MAP estimation

Let the finite set $\mathcal{T}_m \subset \mathcal{T}$ be the time instants where measurements are taken. In addition, let the \mathbb{R}^{m_t} -valued random variables Y_t denote the measurements taken at each $t \in \mathcal{T}_m$, with the dimension of the measurement vector m_t possibly varying across different instants. We then make the following assumptions.

Assumption 6 (measurements).

- Each measurement Y_t is conditionally independent, given the state X_t at the measurement instant, of all the states X_τ and measurements Y_τ at the remaining time instants $\tau \neq t$.
- Each measurement Y_t admits a conditional density $\psi_t(y|x)$, given the state X_t , which is of class C_b^0 with respect to its second argument.
- All measured values y_t have a nonzero predictive prior density, i.e., $\mathbb{E}[\psi_t(y_t|X_t)] > 0$.

Given a sequence y_t of measurement values, we wish to obtain the mode of the state path conditioned on $\mathcal{Y} := \{\omega \in \Omega: Y_t(\omega) = y_t, \forall t \in \mathcal{T}_m\}$. We then have the following result.

Theorem 7. *Let $\phi, \varphi \in \mathcal{H}$ with $\psi_t(y_t | \varphi(t)) > 0$ for all $t \in \mathcal{T}_m$. In addition, let \mathcal{A}_1^ϵ and \mathcal{A}_2^ϵ be two families of events measurable with respect to the σ -algebra generated by X such that $\mathcal{A}_1^\epsilon \subset \mathcal{B}_\phi^\epsilon$, $\mathcal{A}_2^\epsilon \subset \mathcal{B}_\varphi^\epsilon$ and $P(\mathcal{A}_1^\epsilon)/P(\mathcal{A}_2^\epsilon)$ converges as $\epsilon \downarrow 0$. Then, for a system with state process satisfying Assumption 3 and measurements satisfying Assumption 6,*

$$\lim_{\epsilon \downarrow 0} \frac{P(\mathcal{A}_1^\epsilon | \mathcal{Y})}{P(\mathcal{A}_2^\epsilon | \mathcal{Y})} = \frac{\prod_{t \in \mathcal{T}_m} \psi_t(y_t | \phi(t))}{\prod_{t \in \mathcal{T}_m} \psi_t(y_t | \varphi(t))} \lim_{\epsilon \downarrow 0} \frac{P(\mathcal{A}_1^\epsilon)}{P(\mathcal{A}_2^\epsilon)}. \quad (8)$$

Proof. Using Bayes' rule, the posterior probability of any event \mathcal{A} measurable with respect to the σ -algebra generated by X is

$$P(\mathcal{A}|\mathcal{Y}) = \frac{P(\mathcal{A}) \prod_{t \in \mathcal{T}_m} \mathbb{E}[\psi_t(y_t|X_t) | \mathcal{A}]}{\prod_{t \in \mathcal{T}_m} \mathbb{E}[\psi_t(y_t|X_t)]}. \quad (9)$$

Due to the continuity of the measurement likelihood ψ_t with respect to its second argument, we have that for all $t \in \mathcal{T}_m$,

$$\lim_{\epsilon \downarrow 0} \mathbb{E}[\psi_t(y_t|X_t) | \mathcal{A}_1^\epsilon] = \psi_t(y_t | \phi(t)), \quad (10)$$

and analogously for \mathcal{A}_2^ϵ and φ . Substituting (9) and (10) into the left-hand side of (8) we obtain the result on the right-hand side of the equation. \square

The MAP state path estimation problem then follows from Theorem 7, by setting $\mathcal{A}^\epsilon = \mathcal{B}_\phi^\epsilon$. We will use the logarithm of the asymptotic sojourn probability because it is more tractable numerically and simplifies some convergence proofs. As some densities can be equal to zero, the resulting merit functions will take values on the extended real number line $\bar{\mathbb{R}} := \mathbb{R} \cup \{-\infty, \infty\}$.

Lemma 8. *The continuous-discrete MAP state paths are curves $\phi \in \mathcal{H}$ that maximize the MAP merit $H: \mathcal{H} \rightarrow \bar{\mathbb{R}}$ given by*

$$H(\phi) := J(\phi) + \ln \nu(\phi(0)) + \sum_{t \in \mathcal{T}_m} \ln \psi_t(y_t | \phi(t)). \quad (11)$$

If instead $\mathcal{A}^\epsilon = \mathcal{C}_{\eta, x_0}^\epsilon$, Theorem 7 leads to the minimum energy estimation problem.

Lemma 9. *The minimum energy state paths, which are associated with the MAP noise paths, are curves $\phi \in \mathcal{H}$ that maximize the energy merit $H_e: \mathcal{H} \rightarrow \bar{\mathbb{R}}$ given by*

$$H_e(\phi) := J_e(\phi) + \ln \nu(\phi(0)) + \sum_{t \in \mathcal{T}_m} \ln \psi_t(y_t | \phi(t)). \quad (12)$$

We now present the discretization schemes and their corresponding discretized MAP state path estimation problems.

7 Discretized MAP state path estimation

To use discrete-time MAP state path estimation methods in systems with dynamics described by SDEs, the state transition density between points in a partition of the experiment's time interval must be known. As the

transition densities of SDEs do not have closed-form solutions for general nonlinear systems, approximations must be used. In this work, we consider approximations used in the numerical solution of SDEs.

We start with an ordered partition $\mathcal{P} := \{t_k\}_{k=0}^N \subset \mathcal{T}$, which includes the endpoints $t_0 = 0$, $t_N = T$ and measurement times $\mathcal{T}_m \subset \mathcal{P}$. The discretization schemes then represent the evolution of the state at the partition points by difference equations. For brevity when working with partition points, we will use the following notation when unambiguous: X_{t_k} will be shortened to X_k and sequences like X_{t_0}, \dots, X_{t_N} will be shortened to $X_{0:N}$. The sequence $X_{0:N}$ will be denoted the discretized state path.

Change of variables is used to obtain the discretized state path densities. For this we need the well-known theorem below.

Theorem 10 (Georgii [2008, Prop. 9.1]). *Let A be an \mathbb{R}^n -valued random variable admitting a probability density. If $q: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a C^1 diffeomorphism and the \mathbb{R}^n -valued random variable B is defined as $B := q(A)$, then B admits a probability density, which is given by*

$$p_B(b) = |\det \nabla q^{-1}(b)| p_A(q^{-1}(b)), \quad (13)$$

where $\det \nabla q^{-1}(b)$ is the Jacobian determinant of the inverse of q , evaluated at b .

In (13), the Jacobian determinant represents how the diffeomorphism changes the volume element at each point. We now present each discretization and the corresponding discretized MAP state path estimation problems.

7.1 Euler scheme

The stochastic Euler scheme [Kloeden and Platen, 1992, Sec. 9.1], also known as the Euler–Maruyama scheme, is one of the simplest SDE discretization schemes. It approximates the evolution of the states by

$$X_{k+1} = X_k + f(t_k, X_k) \delta_k + G \Delta W_k, \quad (14)$$

where $\delta_k := t_{k+1} - t_k$ and $\Delta W_k := W_{k+1} - W_k$. Recall that the increments of the Wiener process are, by definition, independent and normally distributed with zero mean and covariance Id_k , i.e., $\Delta W_k \sim \mathcal{N}(0, \text{Id}_k)$.

From (14) and the assumptions that f is of class C_b^2 and G is invertible, it can be seen that $X_{0:N}$ can be obtained from $X_0, \Delta W_{0:N-1}$ by a diffeomorphism. Applying Thm. 10, the joint prior density of the state path over the discretization points is given by

$$p(x_{0:N}) = \frac{\nu(x_0) \exp(R(x_{0:N}))}{\prod_{k=0}^{N-1} |\det G| \sqrt{(2\pi)^n \delta_k}}, \quad (15)$$

where R is the Euler energy functional, defined as

$$R(x_{0:N}) := -\frac{1}{2} \sum_{k=0}^{N-1} \delta_k \left\| \frac{\Delta x_k}{\delta_k} - f(t_k, x_k) \right\|_Q^2. \quad (16)$$

Assumption 6a implies that the joint measurement likelihood, i.e., the joint density of all measurements, conditioned on the state path, is given by

$$p(\mathcal{Y}|x_{0:N}) = \prod_{t \in \mathcal{T}_m} \psi_t(y_t | x_t). \quad (17)$$

Applying Bayes' rule we have that, for fixed y_t , the posterior state path density can be obtained from (15) and (17):

$$p(x_{0:N}|\mathcal{Y}) \propto p(\mathcal{Y}|x_{0:N})p(x_{0:N}). \quad (18)$$

Taking the logarithm of (18) and removing the constant terms, which do not influence the location of maxima, we obtain the merit function for MAP estimation.

Lemma 11. *The discretized MAP state paths under the Euler scheme are the vectors $x_{0:N} \in \mathbb{R}^{Nn+n}$ which maximize the Euler merit*

$$S(x_{0:N}) := R(x_{0:N}) + \ln \nu(x_0) + \sum_{t \in \mathcal{T}_m} \ln \psi_t(y_t | x_t). \quad (19)$$

7.2 Trapezoidal scheme

The trapezoidal scheme [Kloeden and Platen, 1992, p. 500] represents the SDE over the partition points by an implicit difference equation:

$$X_{k+1} = X_k + \frac{\delta_k}{2} [f(t_k, X_k) + f(t_{k+1}, X_{k+1})] + G \Delta W_k. \quad (20)$$

Since the drift f is assumed to be of class C_b^2 , for sufficiently small δ_k (20) is a contraction mapping and always has a unique solution. We shall assume that the partition is sufficiently fine such that this holds.

As with the Euler scheme, applying Thm. 10 we obtain the joint discretized state path prior density:

$$p(x_{0:N}) = \frac{\nu(x_0) \exp(U(x_{0:N}))}{\prod_{k=0}^{N-1} |\det G| \sqrt{(2\pi)^n \delta_k}}, \quad (21)$$

where U is the trapezoidal Onsager–Machlup functional defined as

$$U(x_{0:N}) := \sum_{k=0}^{N-1} \ln \det \left(I - \frac{1}{2} \nabla_x f(t_{k+1}, x_{k+1}) \delta_k \right) - \frac{1}{2} \sum_{k=0}^{N-1} \delta_k \left\| \frac{\Delta x_k}{\delta_k} - \frac{1}{2} [f(t_k, x_k) + f(t_{k+1}, x_{k+1})] \right\|_Q^2$$

with $\nabla_x f(t, a)$ being the Jacobian matrix of f with respect to x , evaluated at t, a . Using equations (17), (18) and (21) and taking their logarithm, we obtain the discretized MAP state path estimation problem using the trapezoidal scheme.

Lemma 12. *The discretized MAP state paths under the trapezoidal scheme are the vectors $x_{0:N} \in \mathbb{R}^{Nn+n}$ which maximize the trapezoidal merit*

$$V(x_{0:N}) := U(x_{0:N}) + \ln \nu(x_0) + \sum_{t \in \mathcal{T}_m} \ln \psi_t(y_t | x_t). \quad (22)$$

Unlike the Euler scheme, in the trapezoidal the state X_{k+1} depends nonlinearly on the Wiener process increment ΔW_k , due to the implicit nature of (20). This means that the prior density of a state path does not depend only on the density of the associated noise, but also on the change of the volume element, represented by the Jacobian determinant of the drift function. This extra term represents how much the system dynamics amplifies or attenuates the noise around each path, akin to the drift divergence in the Onsager–Machlup functional. We note that this term has a numerical and conceptual connection to Lyapunov exponents, whose sum quantifies the average change in hypervolume due to perturbations of a path of a noise-free discrete-time map.

We will now present the theory of hypographical convergence to relate the discretized and continuous-discrete MAP estimation problems.

8 Limit of the discretized problems

We now consider the discretized problems over a nested sequence of partitions $\{\mathcal{P}_i\}_{i=1}^\infty$. The last index of each partition will be denoted N_i , their elements by t_{ki} and the time increments by $\delta_{ki} := t_{k+1,i} - t_{ki}$. In addition, $\bar{\delta}_i$ will denote the mesh (or norm) of the partition \mathcal{P}_i , i.e., the largest distance between consecutive points

$$\bar{\delta}_i := \max_{0 \leq k < N_i} \delta_{ki}.$$

We consider partition sequences which satisfy

$$\mathcal{P}_i \subset \mathcal{P}_{i+1}, \quad \lim_{i \rightarrow \infty} \bar{\delta}_i = 0, \quad \bar{\delta}_i N_i < c_3, \quad (23)$$

where $c_3 > 0$ is an upper bound.

To relate the various discretizations and the continuous-discrete problems, we must formulate all problems on the same space. We do this by representing the discretized state paths $\phi_i: \mathcal{T} \rightarrow \mathbb{R}^n$ as the piecewise linear interpolants of given $x_{0:N_i}$. We also denote by $\mathcal{H}_i \subset \mathcal{H}$ the space of all piecewise linear functions from \mathcal{T} to \mathbb{R}^n with breaks along \mathcal{P}_i . The discretized merit functions for each

partition will be denoted by S_i and V_i , and are accordingly extended for inputs from \mathcal{H}_i . The same goes for the Euler energy functional R_i and the trapezoidal Onsager–Machlup functional U_i .

The Cameron–Martin space \mathcal{H} is a reproducing kernel Hilbert space. It can be endowed with the following inner product and its associated norm:

$$\langle a, b \rangle_{\mathcal{H}} := a(0)^{\top} b(0) + \int_0^T \dot{a}(t)^{\top} \dot{b}(t) dt.$$

This norm for \mathcal{H} dominates the supremum norm [Daniel, 1969, Prop. 4.1], i.e., there exists a constant $c_4 \in \mathbb{R}$ such that, for all $\phi \in \mathcal{H}$,

$$\sup_{t \in \mathcal{T}} \|\phi(t)\| \leq c_4 \|\phi\|_{\mathcal{H}}. \quad (24)$$

We use hypographical convergence, or hypo-convergence, to relate the discretized problems and their continuous-discrete counterparts. Hypo-convergence is a powerful tool because when a series of optimization problems hypo-converge, then any cluster point of their solutions solves the limit problem [Polak, 1997, Thm. 3.3.3]. Hypographical limits are formally defined as the limit of the hypographs of a series of optimization problems but, in metric spaces, the following definition is equivalent and more convenient to use.

Theorem 13 (adapted from Polak [1997, Thm. 3.3.2]). *Let \mathcal{H} be a normed linear space and $\{\mathcal{H}_i\}_{i=1}^{\infty}$ be a nested family of finite-dimensional subspaces whose limit \mathcal{H}_{∞} is dense in \mathcal{H} , i.e.,*

$$\mathcal{H}_i \subset \mathcal{H}_{i+1} \subset \mathcal{H}, \quad \mathcal{H}_{\infty} := \bigcup_{i=1}^{\infty} \mathcal{H}_i, \quad \text{clos } \mathcal{H}_{\infty} = \mathcal{H},$$

where clos stands for topological closure. Additionally, let $H: \mathcal{H} \rightarrow \overline{\mathbb{R}}$ be the upper semi-continuous merit function of a maximization problem and $H_i: \mathcal{H}_i \rightarrow \overline{\mathbb{R}}$ be a family of upper semi-continuous merit functions of a family of maximization problems. A sufficient condition for the family of maximization problems to hypo-converge to the maximization of H is that for every convergent infinite sequence $\{\phi_i\}_{i \in \mathcal{I}}$ where $\phi_i \in \mathcal{H}_i$ and $\phi_i \rightarrow \phi$, then

$$\limsup_{i \rightarrow \infty} H_i(\phi_i) = H(\phi).$$

We now present a few lemmas that will help prove hypo-convergence.

Lemma 14. *Let $q: \mathcal{T} \times \mathbb{R}^n \rightarrow \mathbb{R}^o$ be a continuous function and $\{\phi_i\}_{i \in \mathcal{I}}$ be a convergent infinite sequence where $\phi_i \in \mathcal{H}_i$ and $\phi_i \rightarrow \phi$. Additionally, let $g_i: \mathcal{T} \rightarrow \mathbb{R}^n$ be the piecewise constant interpolations of $q(t, \phi_i(t))$ with breaks along the partitions \mathcal{P}_i and define*

$g(t) := q(t, \phi(t))$. Then g_i and g are bounded and

$$\lim_{i \rightarrow \infty} \sup_{t \in \mathcal{T}} \|g(t) - g_i(t)\| = 0.$$

Proof. From (24) and the fact that the ϕ_i sequence is convergent with respect to the \mathcal{H} -norm, it follows that $\sup_{i,t} \|\phi_i(t)\| < \infty$. We can then restrict our analysis of q to a compact subset of its domain, where it is bounded and uniformly continuous. Let ρ_q denote a modulus of continuity of q , i.e., for all $t, \tau \in \mathcal{T}$ and a, b in the compact subset of \mathbb{R}^n being analysed,

$$\|q(t, a) - q(\tau, b)\| \leq \rho_q(|t - \tau|, \|a - b\|). \quad (25)$$

Similarly, let ρ_{ϕ} denote the modulus of continuity of ϕ and define $e_i := c_4 \|\phi_i - \phi\|_{\mathcal{H}} + \rho_{\phi}(\bar{\delta}_i)$. Then $e_i \rightarrow 0$ and

$$\sup_{t, \tau \in [t_{k,i}, t_{k+1,i}]} \|\phi_i(t) - \phi(\tau)\| \leq e_i. \quad (26)$$

Combining (25) and (26) we have that

$$\sup_{t \in \mathcal{T}} \|g_i(t) - g(t)\| \leq \rho_q(\bar{\delta}_i, e_i), \quad \rho_q(\bar{\delta}_i, e_i) \rightarrow 0. \quad \square$$

We now prove that, for any convergent sequence of piecewise linear functions over the partitions, the discretized merits converge to the continuous-discrete merits. The L^2 space for \mathbb{R}^n -valued functions will be denoted by L_n^2 , and corresponds to the direct sum of n copies of the standard L^2 space, with its inner product given by

$$\langle a, b \rangle_{L_n^2} := \int_0^T a(t)^{\top} b(t) dt.$$

Lemma 15. *Let $\{\phi_i\}_{i \in \mathcal{I}}$ be a convergent infinite sequence where $\phi_i \in \mathcal{H}_i$ and $\phi_i \rightarrow \phi$. Then*

$$\lim_{i \rightarrow \infty} S_i(\phi_i) = H_e(\phi),$$

where H_e is defined in (12) and S_i is given by (19) for partition \mathcal{P}_i .

Proof. Since by (24) convergence in \mathcal{H} implies uniform convergence and the initial state density ν and measurement likelihoods ψ_t are continuous, we have that

$$\nu(\phi_i(0)) \rightarrow \nu(\phi(0)), \quad \psi_t(\phi_i(t)) \rightarrow \psi_t(\phi(t)),$$

so, as the logarithm is also continuous in $\overline{\mathbb{R}}$, it only remains for us to prove that $R_i(\phi_i) \rightarrow J_e(\phi)$.

Let g_i be the piecewise constant interpolation of $f(t, \phi_i(t))$ with breaks across \mathcal{P}_i and with the left end-point of each piece as the interpolation point. In addition, define $g(t) := f(t, \phi(t))$. Then, using the bound

$\|x\|_Q \leq \|Q\| \|x\|$ and the fact that $\dot{\phi}_i$ is piecewise constant, from the definitions (7) and (16) of the functionals we have that

$$|R_i(\phi_i) - J_e(\phi)| \leq \|Q\| \left(\|\dot{\phi}_i - g_i\|_{L_n^2}^2 - \|\dot{\phi} - g\|_{L_n^2}^2 \right).$$

As the norm and exponentiation are continuous functions, it suffices to show that $\phi_i \rightarrow \phi$ and $g_i \rightarrow g$ with respect to the L_n^2 norm. The former follows trivially from the fact that $\|\dot{\phi}\|_{L_n^2} \leq \|\varphi\|_{\mathcal{H}}$ for all $\varphi \in \mathcal{H}$. The latter follows from Lem. 14 and the fact that uniform convergence implies convergence in L^2 for finite measure spaces. \square

Lemma 16. *Let $\{\phi_i\}_{i \in \mathcal{I}}$ be a convergent infinite sequence where $\phi_i \in \mathcal{H}_i$ and $\phi_i \rightarrow \phi$. Then*

$$\lim_{i \rightarrow \infty} V_i(\phi_i) = H(\phi),$$

where H is defined in (11) and V_i is given by (22) for partition \mathcal{P}_i .

Proof. As in the proof of Lem. 15, the convergence of the terms corresponding to ν and ψ_t follows trivially, so we will concentrate on proving that $U_i(\phi_i) \rightarrow J(\phi)$. Define

$$e_i := \sum_{k=0}^{N-1} \ln \det \left(\mathbf{I} - \frac{1}{2} \nabla_x f(t_{k+1}, \phi_i(t_{k+1})) \delta_k \right) - \frac{1}{2} \int_0^T \operatorname{div} f(t, \phi(t)) dt.$$

We now show that $e_i \rightarrow 0$.

First, recall that for any nonsingular matrix $A \in \mathbb{R}^{n \times n}$, $\ln \det A = \operatorname{tr} \ln A$, where $\ln A$ is any matrix logarithm of A . In addition, for the principal logarithm we have that, for all matrices A with spectral radius smaller than unity, $\ln(\mathbf{I} + A) = \sum_{j=1}^{\infty} (-1)^{j+1} A^j / j$ [Higham and Al-Mohy, 2010, Sec. 5.2]. Since the drift f is assumed to be of class C_b^2 , we have that for small enough δ_i (or, equivalently, sufficiently large i) $\nabla_x f(t, x) \delta_{ki}$ always has spectral radius smaller than unity. This implies that there exists a constant c_{10} such that, for all i above a certain limit, $x \in \mathbb{R}^n$ and $t \in \mathcal{T}$,

$$|\ln \det(\mathbf{I} - \frac{1}{2} \nabla_x f(t, x) \delta_{ki}) - \frac{1}{2} \operatorname{div} f(t, x) \delta_{ki}| \leq c_{10} \bar{\delta}_i^2.$$

Let d_i denote the piecewise constant interpolation of $\operatorname{div} f(t, \phi_i(t))$ with breaks across \mathcal{P}_i and with the right endpoint of each piece as the interpolation point. Likewise, define $d(t) := \operatorname{div} f(t, \phi(t))$. We then have that, for all sufficiently large i ,

$$|e_i| \leq \frac{1}{2} \int_0^T |d(t) - d_i(t)| dt + c_{10} N_i \bar{\delta}_i^2.$$

From Lem. 14 and the fact that uniform convergence implies convergence in L^1 in finite measure spaces we have that $d_i \rightarrow d$ in L^1 . Also, from (23) we have that $N_i \bar{\delta}_i$ is bounded, so $e_i \rightarrow 0$.

We now show that the quadratic terms of the merit functions converge. Let g_i and h_i denote piecewise constant interpolations of $f(t, \phi_i(t))$ with breaks across \mathcal{P}_i ; g_i with the left endpoint of each piece as the interpolation point and h_i with the right endpoint. It suffices to prove that

$$\lim_{i \rightarrow \infty} \|\dot{\phi}_i - \frac{1}{2}(g_i + h_i)\|_{L_n^2}^2 - \|\dot{\phi} - g\|_{L_n^2}^2 = 0.$$

From Lem. 14 we have that $g_i \rightarrow g$ and $h_i \rightarrow g$ in L_n^2 . Furthermore, $\phi_i \rightarrow \phi$, so due to continuity of the norm and exponentiation $V_i(\phi_i) \rightarrow H(\phi)$. \square

Lemma 17. *The space \mathcal{H}_{∞} is dense in \mathcal{H} .*

Proof. Denote by \mathcal{S}_i the space of piecewise constant functions from \mathcal{T} to \mathbb{R}^n with breaks along \mathcal{P}_i . The \mathcal{S}_i are spaces of order 1 splines and their union is dense in the space of continuous functions equipped with the supremum norm [de Boor, 2001, p. 147]. Since continuous functions are dense in L_n^2 and denseness is a transitive relation, \mathcal{S}_i is dense in L_n^2 . Thus, for any $\phi \in \mathcal{H}$ and $\epsilon \in \mathbb{R}_{\geq 0}$, there exists an $i \in \mathbb{N}$ and $u \in \mathcal{S}_i$ such that $\|u - \dot{\phi}\|_{L_n^2} < \epsilon$. If we then define $\phi_i(t) := \phi(0) + \int_0^t \dot{\phi}(\tau) d\tau$ we have that $\phi_i \in \mathcal{H}_i$ and $\|\phi_i - \phi\|_{\mathcal{H}} < \epsilon$. \square

We are now ready to prove one of the main results of this paper.

Theorem 18. *The discretized MAP state path estimation with the trapezoidal scheme hypo-converges to the continuous-discrete MAP state path estimation; and the discretized MAP state path estimation with the Euler scheme hypo-converges to the minimum energy estimation.*

Proof. From Lem. 17 we know that \mathcal{H}_{∞} is dense in \mathcal{H} . From Lemmas 15 and 16 we then have that the merits converge for any convergent sequence of ϕ_i . \square

In the sequence, we present simulated examples to illustrate an application of robust estimation under the presence of outliers using both the MAP and minimum energy estimators.

9 Simulated Examples

The stochastic Van der Pol oscillator is a nonlinear system commonly used to compare state estimators. This system, discretized with the Euler scheme, was used by Aravkin et al. [2011a, 2012a] for a comparison of

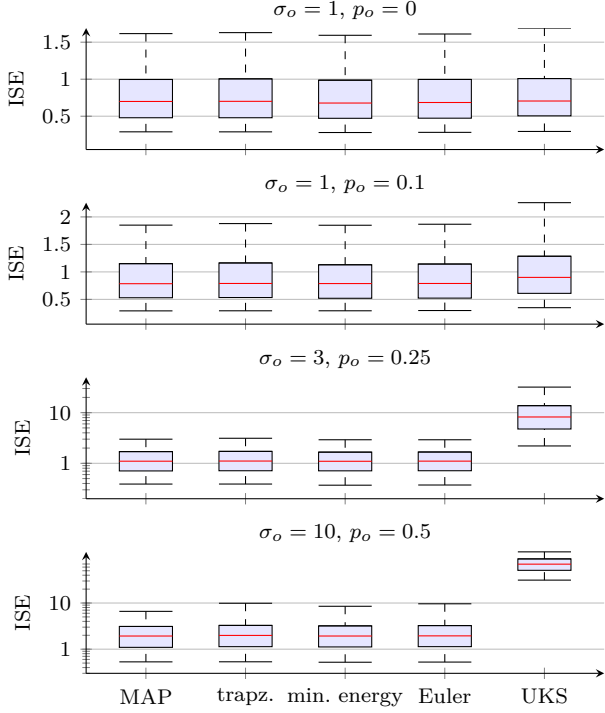


Fig. 3. Box and whisker plot of the integrated squared error for the robust estimation. The whiskers indicate the 5th and 95th percentiles.

robust discrete-time estimators. Denoting its states by $x := [u, v]^T$, its dynamics are defined by

$$f(t, x) := \begin{bmatrix} v \\ -u + 2(1 - u^2)v \end{bmatrix}, \quad G := \begin{bmatrix} 0.1 & 0 \\ 0 & 0.1 \end{bmatrix}.$$

This system has a nonlinear drift divergence, given by $\text{div } f(t, x) = 2(1 - u^2)$, which is higher near the origin. Consequently, the MAP merit will favor paths with larger u while the energy merit will not make such distinctions. To highlight this difference between the estimators, the initial state was chosen around the origin, drawn from the normal distribution $\mathcal{N}([0, 0]^T, \frac{1}{100}\mathbf{I})$.

The system was simulated using the strong explicit order 1.5 scheme [Kloeden and Platen, 1992, Sec. 11.2] with time step 5×10^{-4} and experiment duration $T = 16$. The measurements were generated according to the Gaussian mixture distribution

$$Y_t | X_t \sim (1 - p_o)\mathcal{N}(U_t, \sigma_y^2) + p_o\mathcal{N}(U_t, \sigma_o^2),$$

where $\sigma_y = 0.5$ is the regular measurements' standard deviation, σ_o is the outliers' standard deviation and p_o is the outlier probability, the latter two of which were varied across experiments. The measurements were taken at time intervals of 0.1, i.e., $\mathcal{T}_m = \{0, 0.1, \dots, 16\}$.

Student's t -distribution was used as the measurement likelihood for the MAP, minimum energy and discretized estimators to confer robustness against outliers. The expression of the log-likelihood used in the estimation is

$$\ln \psi_t(y|x) = -\frac{5}{2} \ln \left(1 + \frac{(u-y)^2}{4\sigma_y^2} \right),$$

which corresponds to the t -distribution with 4 degrees of freedom and scaling σ_y . The estimates were also compared against the unscented Kalman smoother (UKS) of Särkkä [2008], using the continuous-discrete SDE prediction step of Arasaratnam et al. [2010]. The measurement covariance of σ_y^2 was used for the UKS.

Optimal control techniques were used to solve the MAP and minimum energy estimation problems. The estimation was transcribed to a nonlinear programming (NLP) problem using a third-order Legendre–Gauss–Lobatto direct collocation method equivalent to the Hermite–Simpson method [Betts, 2010, Sec. 4.5]. The resulting NLP was then solved using the IPOPT solver of Wächter and Biegler [2006], part of the open-source COIN-OR project. The discretized MAP state path estimation problems with the Euler and trapezoidal schemes were also solved using the IPOPT solver. The estimation time step of 5×10^{-3} was used for all estimators.

The experiment and estimation were repeated 2,000 times for each combination of σ_o, p_o and the integrated squared error (ISE) was calculated for each estimate:

$$\text{ISE} = \int_0^T (X_t - x(t))^T (X_t - x(t)) dt,$$

where X is the simulated state and x its estimate. The results are summarized in Fig. 3. It can be seen that, in all scenarios, the performance of the minimum energy and MAP estimators was comparable. However, the performance of the UKS becomes significantly worse as the outliers become more frequent and with a higher variance. These results mirror those of Aravkin et al. [2012a].

It should be noted that although the ISE distribution of both the MAP and minimum energy estimates was similar, the estimates were often different. This point can be seen in Fig. 4, which shows a portion of the simulated and estimated timeseries for one experiment. It is not apparent, however, whether any one is intrinsically better statistically than the other, except in very specific scenarios. *They represent different statistics with different interpretations*; it is up to the user to determine which one is preferable for each application.

We also note that, for systems with student's t -distribution as the measurement likelihood, the discretized MAP state path estimator with the Euler scheme is the t -robust smoother of Aravkin et al.

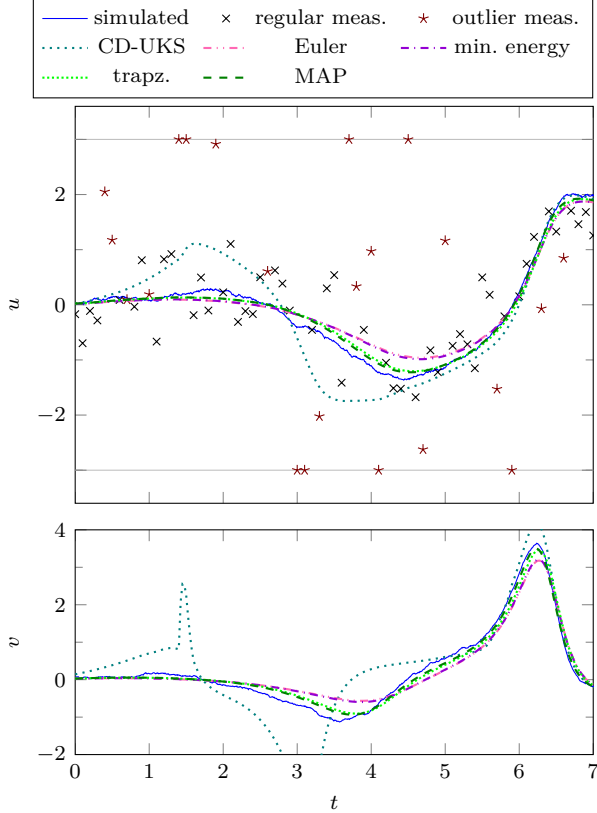


Fig. 4. Example of a portion of the Van der Pol simulated experiment timeseries for $\sigma_o = 3$, $p_o = 0.25$. Outliers outside the plot range are represented at ± 3 . The MAP estimator favors paths with larger u .

[2012a]. Therefore, a practical consequence of the hypoconvergence result of Thm. 18 is that, for sufficiently small discretization steps, the estimates of the t -robust smoother should be close to the minimum energy estimates. This was verified in the simulated experiments, and can be observed in Figs. 3 and 4. The MAP state path estimates were also close to the discretized state path estimates with the trapezoidal discretization.

10 Conclusions and Future Work

In this paper, we addressed MAP state path estimation in continuous-discrete systems described by SDEs. We began with a formal definition of mode for continuous-time stochastic processes so that the target statistic and its interpretation are well-defined. We then showed how the Onsager–Machlup functional can be used to construct the MAP state path and the MAP noise path estimators, the latter of which is associated with the minimum energy state paths. In this way, a clear statistical interpretation of the minimum energy state paths was provided, together with the conditions under which they coincide with the MAP state paths.

Additionally, we showed that the popular technique to

solve these estimation problems by discretizing the dynamics and obtaining the discretized MAP state path estimate leads to different results depending on the discretization scheme used. The Euler scheme, which is the simplest and therefore most widely used method for discretizing nonlinear SDEs, generates discretized estimators which hypo-converge to the MAP *noise* path estimation. The trapezoidal scheme, on the other hand, generates discretized estimators which hypo-converge to the MAP *state* path estimation.

As mentioned in the end of Sec. 7, the difference between the energy and Onsager–Machlup functionals can be interpreted as accounting for the amplification of noise by the system. The latter favours paths around more locally stable regions of the state-space, quantified by the sum of the eigenvalues of the drift Jacobian matrix $\nabla_x f$, i.e., the drift divergence $\text{div } f$. Paths around less stable regions of state space amplify more greatly the perturbations due to noise, which reduces the probability that the state stays inside small tubes around it.

Having proved the link between discretized state path estimation and variational problems also opens the possibility of using optimal control techniques to solve the former, as was done in the simulated examples. These methods are better designed for the solution of infinite-dimensional optimization and are characterized by their order of convergence to the variational solution. SDE discretization methods, on the other hand, are characterized by order of convergence in mean (or of the means) of the simulated and true stochastic processes [Kloeden and Platen, 1992, Secs. 9.6 and 9.7]. There also exist optimal control methods of arbitrarily high order; SDE methods for MAP calculation are limited to strong order 1.5 or weak order 4.

The results herein can be generalized in a number of ways. Rank-deficient G matrices can be allowed by using the same approach as Aihara and Bagchi [1999b], for example. Hypoconvergence can also be used to analyze the limits of discretized joint MAP state path and parameter estimation [Dutra et al., 2012]. Finally, other popular higher order discretization schemes such as the order 1.5 Itô–Taylor should generate problems which hypo-converge to MAP state path estimation as well.

References

- S.I. Aihara and A. Bagchi. On the Mortensen equation for maximum likelihood state estimation. *IEEE Trans. Autom. Control*, 44(10):1955–1961, 1999a.
- S.I. Aihara and A. Bagchi. On maximum likelihood nonlinear filter under discrete-time observations. In *IEEE ACC*, volume 1, pages 450–454, 1999b.
- I. Arasaratnam, S. Haykin, and T.R. Hurd. Cubature Kalman filtering for continuous-discrete systems: Theory and simulations. *IEEE Trans. Signal Process.*, 58(10):4977–4993, 2010.

- A.Y. Aravkin, B.M. Bell, J.V. Burke, and G. Pillonetto. An ℓ_1 -Laplace robust Kalman smoother. *IEEE Trans. Autom. Control*, 56(12):2898–2911, 2011a.
- A.Y. Aravkin, B.M. Bell, J.V. Burke, and G. Pillonetto. Learning using state space kernel machines. In *IFAC*, pages 2296–2302, 2011b.
- A.Y. Aravkin, J.V. Burke, and G. Pillonetto. Robust and trend-following Kalman smoothers using Student’s t . In *SYSID*, pages 1215–1220, 2012a.
- A.Y. Aravkin, J.V. Burke, and G. Pillonetto. A statistical and computational theory for robust and sparse Kalman smoothing. In *SYSID*, pages 894–899, 2012b.
- B.M. Bell. The iterated Kalman smoother as a Gauss–Newton method. *SIAM J. on Optimiz.*, 4(3):626, 1994.
- B.M. Bell, J.V. Burke, and G. Pillonetto. An inequality constrained nonlinear Kalman–Bucy smoother by interior point likelihood maximization. *Automatica*, 45(1):25–33, 2009.
- J.T. Betts. *Practical methods for optimal control and estimation using nonlinear programming*. SIAM, 2nd edition, 2010.
- M. Capitaine. Onsager–Machlup functional for some smooth norms on Wiener space. *Probab. Theory Rel.*, 102(2):189–201, 1995.
- H. Cox. *Estimation of state variables for noisy dynamic systems*. PhD thesis, Mass. Inst. Technol., 1963.
- J.W. Daniel. On the approximate minimization of functionals. *Math. Comput.*, 23:573–581, 1969.
- F.E. Daum. Exact finite-dimensional nonlinear filters. *IEEE Trans. Autom. Control*, 31:616–622, 1986.
- C. de Boor. *A Practical Guide to Splines*. Springer, 2001.
- D.A. Dutra, B.O.S. Teixeira, and L.A. Aguirre. Joint maximum a posteriori smoother for state and parameter estimation in nonlinear dynamical systems. In *SYSID*, pages 900–905, 2012.
- S. Farahmand, G.B. Giannakis, and D. Angelosante. Doubly robust smoothing of dynamical processes via outlier sparsity constraints. *IEEE Trans. Signal Process.*, 59(10):4529–4543, 2011.
- T. Fujita and S.-I. Kotani. The Onsager–Machlup function for diffusion processes. *J. Math. Kyoto U.*, 22(1):115–130, 1982.
- H.-O. Georgii. *Stochastics: Introduction to Probability and Statistics*. Walter de Gruyter, 2008.
- S. Godsill, A. Doucet, and M. West. Maximum a posteriori sequence estimation using Monte-Carlo particle filters. *Ann. Stat. Math.*, 53(1):82–96, 2001.
- N.J. Higham and A.H. Al-Mohy. Computing matrix functions. *Acta Numer.*, 19:159–208, 2010.
- P.E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Springer, 1992.
- A. Monin. Modal trajectory estimation using maximum Gaussian mixture. *IEEE Trans. Autom. Control*, 58(3):763–768, 2013.
- E. Polak. *Optimization: algorithms and consistent approximations*. Springer, 1997.
- C. Robert. *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Springer, 2 edition, 2001.
- S. Särkkä. Unscented Rauch–Tung–Striebel smoother. *IEEE Trans. Autom. Control*, 53(3):845–849, 2008.
- S. Särkkä and A. Solin. On continuous-discrete cubature Kalman filtering. In *SYSID*, pages 1221–1226, 2012.
- A. Wächter and L.T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.*, 106(1):25–57, 2006.
- O. Zeitouni. On the onsager–machlup functional of diffusion processes around non C^2 curves. *Ann. Probab.*, 17:1037–1054, 1989.
- O. Zeitouni and A. Dembo. A maximum a posteriori estimator for trajectories of diffusion processes. *Stochastics*, 20(3):221–246, 1987.